REMARKS

The present response is intended to be fully responsive to all points of objection and/or rejection raised by the Examiner and is believed to place the application in condition for allowance. Favorable reconsideration and allowance of the application is respectfully requested.

Applicants assert that the present invention is new, non-obvious and useful. Prompt consideration and allowance of the claims is respectfully requested.

Status of Claims

Claims **1-14** were pending in the application.

Claim **4** has been cancelled.

Claims **1-3**, and **5-14** remain pending in the application

Claim Objections and 35 U.S.C. § 112 Rejections

In the office action, the Examiner objected to Claim 4 under 37 CFR 1.75(c) as being of improper dependent form for failing to further limit the subject matter of a previous claim. The Examiner also rejected claim 4 under 35 U.S.C. § 112, second paragraph, as being vague and indefinite. In order to accelerate the prosecution of this application, Applicants have cancelled Claim 4. The cancellation of claim 4 is not being made to overcome prior art.

35 U.S.C. § 102 Rejections

In the Office Action, the Examiner rejected claims **1, 4, 6-8, and 10-14** under 35 U.S.C. § 102(b), as being anticipated by Jacobson et al (US 5,696,934, hereinafter "Jacobson").

The Examiner has also rejected claims **1, 4, 6, 13-14** under 35 U.S.C. § 102(b), as being anticipated by Craig (US 5,790,176, hereinafter "Craig").

Applicants respectfully traverse the rejection of **1, 4, 6-8, and 10-14** under 35 U.S.C. § 102(b), as being anticipated by Jacobson, and respectfully traverse the rejection of **1, 4, 6, 13-14** under 35 U.S.C. § 102(b), as being anticipated by Craig.

Applicants respectfully assert that neither Jacobson nor Craig, teaches or suggests a method of storage management including storing data on a high reliability high performance storage medium, backing up the data on a high reliability low performance storage medium, and after backing up, copying at least some of the data to a low reliability high performance storage medium, as recited in independent claim 1 of the current application.

In addition, Applicants respectfully assert that neither Jacobson nor Craig, teaches or suggests a storage system including a high performance high reliability storage medium configured for initial storage of data; a low performance high reliability storage medium configured for backup of data initially stored on the high performance high reliability storage medium; and a high performance low reliability storage medium configured to receive at least some of the data after the data has been backed up on the low performance high reliability storage medium, as recited in independent claim 6 of the current application.

Additionally, Applicants respectfully assert that neither Jacobson nor Craig, teaches or suggests a program storage device tangibly embodying a program of instructions to perform method steps of storage management including: storing data on a

high reliability high performance storage medium, backing up the data on a high reliability low performance storage medium, and after backing up, copying at least some of the data to a low reliability high performance storage medium, as recited in independent claim 13 of the current application.

Also, Applicants respectfully assert that neither Jacobson nor Craig, teaches or suggests a computer program product comprising a computer useable medium having computer readable code embodied therein including code for causing the computer to store data on a high reliability high performance storage medium, to back up the data on a high reliability low performance storage medium, and to copy at least some of the data to a low reliability high performance storage medium after the backing up, as recited in independent claim 14 of the current application.

In each of the independent claims 1, 6, 13, and 14 of the current application, paraphrased above, a storage sequence is recited which includes data first being stored on a high reliability high performance storage medium, the data then being backed up on a high reliability low performance storage medium, and then at least some of the data being stored on a low reliability high performance storage medium. The independent claims 1, 6, 13, and 14 clearly indicate that a storage sequence is being recited.

For example, in claim 1, the words "backing up said data" indicate that the data which was stored on the high reliability high performance storage medium is then stored (i.e. backed up) on the high reliability low performance storage medium. Continuing with the example, in claim 1, the words "after said backing up" indicate that the storage of at least some of the data on the low reliability high performance storage medium occurs after the backing up onto the high reliability low performance storage.

As another example, in claim 6 the words "for backup of data initially stored" indicate that the low performance high reliability storage medium is configured to backup

data which was first stored on the high performance high reliability storage medium. Continuing with this example, in claim 6, the words "after said data has been backed up" indicate that the high performance low reliability storage medium is configured to receive data after backup has occurred on the low performance high reliability storage medium. Claims 13 and 14 similarly use terms such as "back(ing) up said data and "after said backing up" to clearly recite the storage sequence.

The storage sequence recited in each of independent claims 1, 6, 13, and 14 of the current application is neither indicated nor hinted at in either of Jacobson or Craig, as will be shown below.

Jacobson describes methods of using storage disks of differing capacity in a single storage volume of a hierarchical disk array. More details are provided in the abstract:

> The methods concern fully utilizing storage capacity in a heterogenous hierarchic disk array having storage disks of differing capacities. The disks are segmented into multiple regions. One method links non-contiguous regions from individual storage disks to form RAID areas. The RAID areas are mapped into a virtual storage space that provides a view of the physical storage space as a single storage volume. Data is then stored in these RAID areas according to different redundancy criteria, such as RAID Level 1 and RAID Level 5. A second method fully utilizes of storage capacity by configuring the heterogeneous disk array to employ a minimum of two equal-sized storage disks that have larger capacity than other individual storage disks in the disk array. The contiguous regions across the multiple disks are then grouped together to form the RAID areas.

In the discussion below, in order to learn from Jacobson what Jacobson is teaching, it is assumed that the categorization in Jacobson of the reliability levels of the various media is correct, without entering into whether the categorization is identical to that of the current application.

In Column 4, lines 19-25 of Jacobson, reliability storage levels are discussed:

> RAID management system 16 provides a data manager means for controlling disk storage and reliability levels, and for transferring data among various reliability storage levels. These reliability storage levels are preferably mirror or parity redundancy levels as described below, but can also include a reliability storage level with no redundancy at all.

In Column 5, lines 28 to 34 of Jacobson, the high reliability of mirror storage is described:

> Mirror group 18 represents a first memory location or RAID area of the disk array which stores data according to a first or mirror redundancy level. This mirror redundancy level is also considered a RAID Level 1. RAID Level 1, or disk mirroring, offers the <u>highest data reliability</u> by providing one-to-one protection in that every bit of data is duplicated and stored within the data storage system. (emphasis added)

In Column 6, lines 31-32 of Jacobson, parity storage is noted as being a low reliability low performance storage compared to mirror storage:

> Parity storage is less expensive than mirror storage, but is also <u>less reliable and has a lower performance.</u> (emphasis added)

Migration of data between parity and mirror storages is discussed in Column 6 line 48 to column 7 line 11 and in column 10, lines 28- to 41, both presented here:

> Data storage system 10 manages the "migration" of data between mirror and parity storage schemes. The management of both types of redundancy is coordinated by RAID management system 16 (FIG. 1). RAID management system 16 manages the two different types of RAID areas in the disk array as a memory hierarchy with the mirror RAID areas acting similar to a cache for the parity RAID areas. Once data is moved from a parity RAID area to a mirror RAID area, the space it once occupied in the parity RAID area is available for storage of other data. RAID management system 16 shifts, organizes, and otherwise manages the data between the mirror and parity RAID areas in

> accordance with a defined performance protocol. <u>The process of moving data between the mirror and parity RAID areas is referred to as "migration".</u>
>
> Data storage system 10 places the more critical data in the mirror RAID areas since this affords the highest performance and reliability. The performance protocols implemented by RAID management system 16 includes one of two preferred migration policies. According to the first migration policy, known as "access frequency", the most frequently accessed data on the hierarchic disk array is maintained in the mirror RAID area 18. Less frequently accessed data is maintained in the parity RAID area 22. According to a second migration policy, known as "access recency", the most recently retrieved data is maintained in the mirror RAID area 18 while the less recently accessed data is stored in parity RAID area 22. Other performance protocols may be employed. Ideally, such protocols are defined based upon the specific computer application and the needs of the user.
>
> As storage capacity approaches full usage, data can be migrated to RAID Level 5 areas to optimize efficient use of the storage space. The proportion of data stored according to RAID Level 5 in relation to data stored according to RAID Level 1 thereby increases. RAID Level 5 uses less space for redundancy data. If all data were stored according to RAID Level 1, half of the storage space would be used for holding redundant data. For RAID Level 5 storage in a disk array of many storage disks, a much smaller fraction of the storage space is dedicated to redundant data. Accordingly, RAID Level 5 permits storage of more user data for a given disk array, but at a cost of lower performance and availability characteristics. Eventually, as storage capacity is filled, all data might be migrated to RAID Level 5 areas. (emphasis added)

As can be seen from the above quotes and from elsewhere in Jacobson, data may be migrated between mirror storage and parity storage. In fact, Jacobson defines migration as "the process of moving data between the mirror and parity RAID areas". Based on Jacobson's reliability characterization of mirror storage and parity storage, any data migration taught by Jacobson is a migration between a high reliability high performance storage (mirror storage) and a low reliability low performance storage (parity storage), as categorized by Jacobson. There is no mention in Jacobson of data migration to a low reliability high performance storage medium nor of data migration to a high reliability low performance storage medium. For this reason, inter-alia, the data

migration discussed in Jacobson does not anticipate the storage sequence recited in the independent claims of the current invention.

Although in Jacobson column 4, lines 24-25 quoted above there is a mention of an optional reliability storage level with no redundancy at all, this no redundancy reliability level is not mentioned elsewhere in Jacobson. It may be speculated that just as Jacobson categorizes the parity storage as being of lower reliability and lower performance than the mirror storage, the no-redundancy reliability level would be categorized by Jacobson as being of lower reliability and lower performance than the parity storage. In any case, what is clear and not subject to speculation is that there is no suggestion in Jacobson that Jacobson categorizes the no redundancy reliability level as being a low reliability high performance storage medium or a high reliability low performance storage medium.

To further understood the teachings of Jacobson, refer to column 4, lines 48 to 65 which describes the memory map store which stores virtual mapping information used to map different storage space onto one another. It should be noted that the memory map store, embodied for example on non-volatile RAM, does <u>not</u> store the data referred to above as being stored in mirror and/or parity storage.

> Data storage system 10 includes a memory map store 21 that provides for <u>persistent storage of the virtual mapping information used to map different storage spaces onto one another.</u> The memory map store is external to the disk array, and preferably resident in the disk array controller 14. The memory mapping information can be continually or periodically updated by the controller or RAID management system as the various mapping configurations among the different views change.
> Preferably, memory map store 21 is <u>embodied as two non-volatile RAMs</u> (Random Access Memory) 21a and 21b which are located in respective controllers 14a and 14b. An example non-volatile RAM (NVRAM) is a battery-backed RAM. A battery-backed RAM uses energy from an independent battery source to maintain the data in the memory for a period of time in the event of power loss to data storage system 10. One preferred construction is a self-refreshing, battery-backed DRAM (Dynamic RAM). (emphasis added).

There is no mention in Jacobson of the data stored in mirror and/or parity storage being migrated to or from the non-volatile RAM, and in fact as mentioned above, migration is defined by Jacobson "as moving data between the mirror and parity RAID areas". It is clear from the quote that the non-volatile RAM is instead used for storing mapping information. Therefore the non-volatile RAM and its explicit or implicit characteristics are irrelevant when comparing the migration of data in Jacobson to the storage sequence recited in the independent claims of the current invention.

In summary, there is neither a hint nor in Jacobson of data being "migrated" to a low reliability high performance storage medium or to a high reliability low performance storage medium. For this reason, inter-alia Jacobson does not anticipate the storage sequence recited in independent claims 1, 6, 13, and 14 of the current invention.

Craig describes a media server for supplying video and multi-media data over the public-switched telephone network. The Media Server includes four levels of storage as noted in the abstract:

> The Media Server includes four levels of storage including DRAM, optical and magnetic disk storage, high speed tape storage and archival storage.

The storage sequence as recited in the independent claims of the current invention and as described above will now be contrasted to Craig. In the discussion of Fig. 5, Craig describes data striping in which data is stored on multiple disks in parallel. The following storage arrangement is stated in Craig in Column 16, lines 7-18:

> Two arrangements for carrying out data striping are illustrated in FIG. 5. The first arrangement is on the left hand side of the controller cluster including 541, 542,

> 543 . . . 54N. And the other arrangement is located on the right hand side. For purposes of this discussion, elements 501, 502, 503 . . . 50N can be considered archival, or slow tapes. Elements 511, 512, 513 . . . 51N can be considered high speed tapes, while elements 521, 522, 523 . . . 52N can be considered disk drives. Finally, elements 531, 532, 533 . . . 53N are DRAM memory elements, generally used as memory caches. Each horizontal line of memory devices is controlled by a CPU such as 541, 542, 543 . . . 54N.

In the discussion below, in order to learn from Craig what Craig is teaching, it is assumed that the categorization in Craig of the performance for the various components of Fig. 5 is correct, without entering into whether the categorization is identical to that of the current application

It will now be examined the various data transfers mentioned between the components of Fig. 5 of Craig. Refer first to Column 17, lines 19 to 32 of Craig:

> The resulting data stream is constituted using a data striping method in which portions of a single program contained in the first type of memory device, such as 501, are distributed over a plurality of different memory devices 511, 512, 513 . . . 51N of a second memory type (preferably having higher operating speed than the first memory type). As depicted in FIG. 5, <u>portions of single programs contained on one, some or all of the memory devices 501, 502, 503 . . . 50N are apportioned among a plurality of memory devices 511, 512, 513 . . . 51N of the second type. For illustrative purposes, the first type of memory device can be considered an archival or slow speed tape while the second type of memory device can be constituted by high speed tape drives.</u> (emphasis added)

It is clear from the above quote from Craig that the described transfer of data is from a slow speed tape (low performance storage medium) to a high speed tape (high performance storage medium).

In Column 17, lines 42 to 61 of Craig, a subsequent stage of transfer of data between the various components of Fig. 5 is described:

> For example, the program portion transmitted between memory devices 501 and 511 via data stream 1(a) can be repeated in a plurality of memory devices by transmitting the same segment from memory device 511 to memory devices 521, 522 and 523 via data streams 11(a), 11(b) and 11(c), respectively. Such repetition is optional depending upon program size and the emphasis placed by the system operator on reliability or the popularity of a particular program. Repetition need not be carried out on a one-to-one basis between memory devices 511, 512, 513 . . . 51N and 521, 522, 523 . . . 52N. Rather, the program segments may simply be further redistributed between the second type of memory device (511 . . . 51N) and the third type of memory device (521 . . . 52N). For purposes of the example depicted in FIG. 5, the third type of memory device can be a magnetic disk drive. Because of the speed of such drives (in comparison to that of high speed tapes) repetition of program segments is more feasible Ln terms of timing efficiency than would be possible with tape drives. (emphasis added)

In the above quote, Craig makes it clear that the subsequent stage includes the transfer of data from the high speed tape to a magnetic disks drive which is even faster than the high speed tape. Therefore again Craig describes a transfer from a lower performance storage to a higher performance storage.

In Column 17, line 62 to Column 18, line 5 of Craig, a subsequent stage of transfer of data between the various components of Fig. 5 is described:

> The fourth type of memory device, depicted by 531, 532, 533 . . . 53N, is a DRAM cache memory. This is the fastest type of memory available and also the most costly for storing large amounts of information. The contents of the third type of memory device represented by 521 . . . 52N, can be fed directly into corresponding memory devices of the fourth type (if they are capable of storing the entirety of the program material in a single memory device such as 521, or, the data from memory device 521 can be distributed in segment at fashion over a plurality of the fast DRAM devices represented by 531 . . . 53N). (emphasis added)

From the above quote in Craig, it is again clear that this stage of transfer of data is to a faster (higher performance) type of memory. Craig calls the target memory, DRAM, "the fastest type of memory available". Therefore again this stage of transfer includes transferring data from a lower performance storage to a higher performance storage.

In an alternative arrangement, Craig in Column 18 lines 32 to 37 mentions that optical disk drives can be substituted for the magnetic disk drives, or for the high speed tape drives:

> In another alternative, optical disk drives such as those depicted by 561, 562, 563 . . . 56N, could be substituted for the magnetic disk drives 521 . . . 52N. If particularly fast magnetic disk drives are being used, the optical disk drives could be substituted for the high speed tape drives 511, 512, 513 . . . 51N.

In the above quote Craig makes it clear that optical disk drives may be substituted for high speed tape drives if the magnetic disk drives are particularly fast. Evidently this condition is necessary so that the transfer of data in this arrangement is also from a lower performance storage to a higher performance (faster) storage.

Therefore, with respect to Figure 5 of Craig, there are one or more possible transfers of data between components. Each transfer of data is from a lower performance storage to a higher performance storage. Therefore the one or more transfers described in Craig do not remotely anticipate the storage sequence as recited in the independent claims of the current invention.

It should also be noted that nowhere in Craig are reliability levels of the various storage components discussed (only levels of performance are mentioned). This also strongly indicates that Craig does not remotely anticipate the storage sequence of the independent claims of the current invention, where the reliability of the various storage media is recited.

In summary Craig does not indicate nor suggest the storage sequence of the current invention as recited in the independent claims.

Based on the above discussion, Applicants respectfully assert that independent claims **1**, **6**, **13**, and **14** are allowable over both Jacobson and Craig. Claims **2**, **3**, **5**, **7-12** depend from, directly or indirectly, claims **1, 6, 13**, and **14** and therefore include all the elements of those claims. Therefore, Applicants respectfully assert that claims **2-3, 5, 7-12** are likewise allowable. Accordingly, Applicants respectfully request that the Examiner withdraw the rejections to independent claims **1, 6, 13** and **14**, and to claims **2-4, 5, 7-12** dependent thereon.

Applicants believe the remarks presented hereinabove to be fully responsive to all of the grounds of rejection raised by the Examiner. In view of these remarks, Applicants respectfully submit that the specification and all of the claims in the present application are in order for allowance. Notice to this effect is hereby requested.

Please charge any fees associated with this paper to deposit account No. 09-0468.

Respectfully submitted,

By: /Stephen C. Kaufman/
Stephen C. Kaufman
Reg. No. 29,551
Phone No. (914) 945-3197

Date: August 14, 2006

IBM Corporation
Intellectual Property Law Dept.
P. O. Box 218
Yorktown Heights, New York 10598